Biologically Inspired Sensor Fusion

26 May 1999

Richard Akita, Robert Pap, et al
Accurate Automation Corp
Chattanooga, TN
& Dr. Joel Davis
Office of Naval Research
Washington, DC

ABSTRACT

Accurate Automation Corporation (AAC) has developed a novel neural network-based sensor fusion system, inspired by the architecture of biological sensor fusion systems. The project performed research into the nature by which information from multiple sensors is fused by the central nervous system and developed a biological model for the process. Based upon this model we developed a system which fuses two or more sensor signals to generate a fused signal with an improved confidence of target existence and position. The system includes gain, control and fusion units, and also include an integration unit. The integration unit receives signals generated by two or more sensors, and generates integrated signals based on the sensor signals. The integration unit performs temporal and weighted spatial integration of the sensor signals, to generate respective sets of integrated signals supplied to the gain control and fusion units. The gain control unit uses a preprogrammed function to map the integrated signals to an output signal that is scaled to generate a gain signal supplied to the fusion unit. The fusion unit uses a preprogrammed function to map its received integrated signals and the gain signal, to a fused signal that is the output of the system. The weighted spatial integration increases the fused signal's sensitivity to near detections and suppresses response to detections relatively distant in space and time, from a detection of interest. The gain control and fusion functions likewise suppress the fused signal's response to low-level signals, but enhances response to high-level signals. In addition, the gain signal is generated from signals integrated over broad limits so that, if a detection occurred near in space or time to a detection of interest, the gain signal will cause the fused signal to be more sensitive to the level of the detection of interest. To our knowledge, the sensor fusion model developed under this effort is the only transition of superior colliculus-based sensor fusion into a functional capability for the Department of the Navy. (Patent Number 5,850,625, issued 12/15/98.)

1. Objective of Sensor Fusion Research

The objective of this research effort was to develop a sensor fusion subsystem that was biologically based which support the Navy's requirements in data fusion. The Department of the Navy's Science and Technology Requirements Guidance (STRG) Aug 98 articulated a number of requirements in sea and area control, command and control, and power projection that requires a reliable and robust data fusion system.

2. Biological Sensor Fusion Phenomenon

For over two decades, there has been a research program in biological sensor data fusion. Research has shown that animals perform fusion from different sensors. This phenomenon has been studied in reptiles, rats, hamsters, ferrets, cats, monkeys and humans. All mammals seem to possess biological circuitry in portions of their brains to perform data fusion among the various stimuli, namely vision, auditory, and somatosensory and that fusion is performed at a neuron level. Researchers at the Bowman-Gray School of Medicine, Wake Forest University have been focusing on understanding the processes of integration (data fusion) of more than one sensory input and more

Form SF298 Citation Data

Report Date *("DD MON YYYY")* 26051999	Report Type N/A	Dates Covered (from... to) *("DD MON YYYY")*

Title and Subtitle Biologically Inspired Sensor Fusion	**Contract or Grant Number**
	Program Element Number
Authors Akita, Richard; Pap, Robert; Davis, Joel	**Project Number**
	Task Number
	Work Unit Number
Performing Organization Name(s) and Address(es) Accurate Automation Corp Chattanooga, TN	**Performing Organization Number(s)**
Sponsoring/Monitoring Agency Name(s) and Address(es)	**Monitoring Agency Acronym**
	Monitoring Agency Report Number(s)
Distribution/Availability Statement Approved for public release, distribution unlimited	
Supplementary Notes	
Abstract	
Subject Terms	
Document Classification unclassified	**Classification of SF298** unclassified
Classification of Abstract unclassified	**Limitation of Abstract** unlimited
Number of Pages 8	

recently, the role that context plays in that set of processes. The portion of the brain involved in attention and orientation behavior is the Superior colliculus (SC). Figure 1 is a model of the structures in the brain. The superior colliculus does not work alone in performing sensor fusion. It acts in concert with inputs from different areas of a set of brain structures called the basal ganglia. The valuable aspect of the basal ganglia is that it provides a form of indirect input into the superior colliculus, which is in addition to the direct input provided by the sensor streams. This is shown in the figure 1.

Inputs from auditory and visual sources each going through low-level processing in the lateral suprasylvian cortex (LS), and from there having direct connections to the superior colliculus (SC). An indirect pathway also exists; the same stimulus passes through the striate cortex (ST) and from there through the substantia nigra pars reticulata (SNR) as well as other cortical areas, and from there to the SC. The SC directly affects motor control. The foundations for this biological sensor fusion architecture are described in [Meredith & Stein, 1986a,b, 1987; Stein et al., 1993; Stein & Meredith, 1990, 1993; Welch & Warren, 1986].

Researchers performed numerous experiments on unimodal, and multisensor fusion. Figure 2 is a representative diagram of the results of these experiments and shows the visual and auditory receptor fields and typical results of unimodal visual, and auditory responses and fused responses. The diagram shows the mapping of the visual and auditory receptive fields for a typical neuron which performs sensor fusion. Several important results can be drawn from the experimental data:

a. When the two stimuli (i.e. visual and auditory) lie within the receptor fields, there is a significant nonlinear response enhancement (above the sum of each of the individual responses, some times as much as 1200% over the unimodal response).
b. When one or more stimuli is outside the receptor field(s), the response from the neuron is significantly depressed.
c. The temporal window during which integration (sensor fusion) can take place varies from cell to cell, but typically extends for 100 – 300 ms. The best multisensor response is most often when latencies for each sensor occur at roughly the same time.
d. When sensory events from different channels are separated in time by a certain amount, the response of the neuron can be depressed.
e. There is a sizeable intersensory overlap; one explanation is that the sensors perform analysis under dynamic conditions where the sensors (the head and eyes) are free to move independently from one another.
f. Although the receptive fields of these sensors can be as large as a hemisphere, animals typically can acquire the location of targets to within fractions of a degree. The ability to acquire targets with a high degree of precision is due to the patterns of convergence that sensors share with the motor structures that control the position of the eyes and head.

Another important result is that for weak signals, multisensory response enhancements were noted, thus we can see that these sensors have a lower detection threshold.

3.0 Model

The architecture for this system involves five different layers, each of which corresponds (conceptually) to a region in the mammalian brain. These include:

a. Two sensor input layers (S1 and S2), The *selective network response* layer (SNR), corresponding to *the substantia nigra pars reticulata* (and associated neural bodies),

b. The first *sensor correlation* layer (SC1), corresponding to the lower levels of the *superior colliculus*, and

c. The second sensor correlation layer (SC2) (optional), which allows for position refinement, and which could relate to the higher levels of the SC.

Figure 3 shows S1 and S2 in the center, with the two SC layers above, and the SNR layer below. A 1-D representation is used for clarity, and was used during development to study parameter effects. Extension to 2-D and 3-D is straightforward, and is covered in our patent. The “bins” in each layer correspond to units of space that are topographically mapped with regard to each other. The figure illustrates a situation where the inputs from a single target are mapped into neighboring locations (respectively) in the S1 and S2 topographic maps. This represents a situation where the target positions as determined by the two sensors are slightly offset from each other. This can be due to either small bias errors between the two sensors, or to different times of observation of a moving target. Fusion is made by combining activations from multiple sensor inputs in two specialized fusion networks, i.e., the *selective network response* layer (SNR) and the *sensor correlation* layer (SC1). (Note that the activation of neurons in the SNR is used to provide a gain control to neurons in SC1.), processing layers of the system. Fusion takes place by correlating detections from corresponding grids.

The following discussion details the sensor fusion system. Figure 4 shows the general block diagram of the sensor fusion system with 2 or more sensors as inputs. For simplicity, we will restrict our discussions for 2 sensors. In its preferred configuration, the system includes an integration unit coupled to receive the first and second sensor signals, and coupled to the gain control and fusion units. The integration unit generates integrated signals based on the first and second sensor signals, and supplies the integrated signals to the gain control and fusion units. The integration unit performs temporal and spatial integration of the sensor signals. The integration unit performs temporal and spatial integration of the two sensor signals, separately and independently. The temporal integration is envisioned to be a graded decay of signal strength. The spatial integration is performed after the temporal integration, with a weighting derived by a difference-of-guassians function, for example. The weighting used by the spatial integration unit reduces noise and conditions the first and second sensor signals to suppress response to detections at positions relatively far from a position currently under analysis, thus improving the determination of target existence at the currently-analyzed position. Also, the integration limits of the first and second sensor signals used to generate the signals output from the integration unit to the gain control unit, are preferably larger in both the temporal and spatial domains than the integration limits of the first and second sensor signals used to generate the signals supplied from the integration unit to the fusion unit. Further, the integration unit optionally uses a predetermined transfer function, preferably a range normalization or scaled sigmoid function, to generate the integrated signals.

The gain control unit is configured to combine the signals received from the integration unit, by inputting the integrated signals into a gain control function that is preprogrammed into the gain control unit. The gain control function maps the levels of the integrated signals to a corresponding signal level that is output from the gain control function. The gain control function is normally set such that it enhances strong detections and suppresses weak detections in the integrated signals. The gain control unit further includes a multiplier that is coupled to receive the signal output from the gain control function, and that multiplies that the signal with a predetermined gain factor, to generate the gain signal that is the output of the gain control unit.

In an alternative configuration, the gain control unit combines the integrated signals input to the gain control unit by averaging the two signals, for example. The gain control unit inputs the averaged signal's level to a gain control function (which in this case is a function of a single variable) that maps the level of the averaged signal to a corresponding level that is output from the gain control function. The output of the gain control function is supplied to the multiplier that multiplies the function's output signal by the gain factor to generate the gain signal.

The fusion unit maps the levels of the signals received from the integration unit, and the level of the gain signal from the gain control unit, to a corresponding level for the fused signal output from the system. More specifically, the fusion unit generates the fused signal based on a predetermined fusion function using the signals from the integration unit and the gain control unit. Similarly to the gain control function, the fusion function generates the fused signal so as to enhance strong detections and suppress weak detections. The fused signal is the ultimate output of the system. Alternatively, the fused signal can be subjected to a final weighting similar to that used in the spatial integration, that is configured to sharpen the location of detections in the fused sensor signal.

The system can also include the first and second sensors that sense the environment and that generate respective first and second sensor signals supplied to the integration unit.

A gain signal is generated based on first and second sensor signals. The method also includes generating a fused signal, based on the first and second sensor and gain signals. The gain and fused signals can be generated with respective predetermined gain control and fusion functions to eliminate noise and sharpen the fused signal to improve the confidence of its target information. The system can also include integrating the first and second sensor signals for use in generating the gain signal and the fused signal. Preferably, the spatial and temporal limits of integration for the first and second sensor signals used to generate the integrated signals supplied to generate the gain signal, are larger than those of the integrated signals supplied to generate the fused signal. Therefore, the integrated signals used to generate the gain signal can be influenced by detections more distant in space or time as compared to the integrated signals used to generated the fused signal. In addition, the method can include a step of weighting that is performed for at least one of the first and second signals during the performance of the steps of integrating the first and second sensor signals. The weighting of the signal(s) can be performed with weight levels derived from a difference-of-gaussians or other similar function. The method can also include the use of a transfer function; for example, a range normalization or scaled sigmoid function, to generate at least one of the integrated first and second signals used in the performance of the integration steps.

4. Results of Concept Testing

We developed a simulation of the sensor fusion system, and ran a suite of experimental trials to determine model performance with different parameter settings. As a summary of the experimental studies, we can say the following:

a. The model allows fusion of target signals separated by small distances, with the probability of fusion initially proportional to the closeness of the two signals. When the target signals overlap or are immediately adjacent, the fusion is 100%, even when the signal-to-noise ratio is small.

As the noise increases, there are multiple effects:

> For small amounts of noise (4<signal-to-noise, S/N<6), targets have a very high probability of being fused when they are close, but this drops off with increase in separation between targets.
>
> As noise increases to a moderate level (2<S/N<4), the probability of fusion decreases – even when the targets are close together. However, there is still reasonable probability of obtaining a fused response even when the targets move apart. The number of separate detections of each of the target presentations and false alarms starts to rise.
>
> As noise becomes severe (S/N<2), the probability of target detection remains low. What is significant is the extent to which the targets are extracted at all within the high noise environment – visual examination of the presentation data will show that it is relatively challenging to discern the targets from noise. This is where the model, especially applied iteratively over time, will have its greatest benefit.

5. Hardware Development: The Sensor Fusion Processor

We designed a unique hardware configuration to implement the sensor fusion method in special-purpose advanced computational hardware. This system, AAC's Sensor Fusion Processor (SFP™) was designed to include the following components: a. AAC's Neural Network Processor (NNP®), b. Input / Output Processor (IOP™), and c. PCI Interface. These three components are organized as is shown in figure 3. The PCI Interface will allow a host computer that has a PCI interface slot to communicate with the SFP™ (specifically to the NNP® component). The PCI Interface would be hosted on a board that will also host the NNP® and the IOP™ (as a daughterboard to the NNP™). It would meet the PCI interface standards that are now common and accepted. The Input-Output Processor (IOP™) would allow the SFP™ to communicate with external devices, providing direct inputs into the SFP™. The IOP™ would provide high-speed I/O directly to the NNP® memory. It would also perform preprocessing of the data. Because it would contain an embedded microprocessor, it could perform weight training. This would make it possible for the IOP™ to adapt weight values for the weights used in the NNP®. The resulting hardware configuration would be suitable for extensions of the basic sensor fusion algorithm, allowing next-generation processing to be hosted on this platform. Finally, because the microprocessor within the IOP™ would eliminate the need for an external CPU hosted in a separate computer, the SFP™ could be used as a stand-alone system, inserted directly into various sensor processing environments. The small size of the SFP™ would make it useful in a wide variety of fieldable applications.

The following subsections describe AAC's design and development efforts, as well as functional capabilities, for each of these components.

5.1 AAC's Neural Network Processor (NNP®)

Figure 5 shows the new dual processor NNP® that makes use of the ISA interface. It can take up to four daughterboards on each side, operating at total capacity with ten processors working simultaneously. This new dual processor design has an improved density of processing components. In addition, issues involving clock loading and distribution as well as power control have been addressed, leading to improved performance.

Each NNP® system is capable of addressing 16K neurons, and contains four transfer function memories. The neural network baseboard can accept different numbers of NNP® modules that allow for tremendous computational diversity. For example, an eight processor NNP® is capable of over one billion connections per second (at 35 MHz). AAC plans to use this large-scale configuration to implement neural network algorithms at high speeds for high-data tasks.
The NNP® architecture is designed to directly implement high level neural network constructs, connections, transfer functions, and levels in a single clock cycle. The NNP® architecture is extremely flexible, permitting any neuron to be interconnected with any other neuron and facilitating the implementation of a wide variety of neural network paradigms. The NNP® supports two I/O buses, one for interprocessor communication, the other for mapping NNP® memory into the memory map of a supporting host processor chip. NNP® software development is supported by an Assembler, an interactive Debugger, and the AAC Neural Network Processor Command Line Tools. We implemented the sensor fusion method on the NNP® and performed algorithm verification. To do the implementation, we wrote the sensor fusion method using the reduced instruction set that is currently in use for the NNP®. Time constraints did not permit us to do a redesign of the RISC instruction set, adding in new capabilities for branching and looping. Thus, specification of a convolution method is a lengthy task. We tested algorithm performance against the outputs of our initial simulations. We found that the transfer function outputs did not match, because of the register size used on the NNP®. In order to configure a system that will demonstrate the SFP™ capabilities, we are installing an NNP® in a PC using an ISA interface. We have checked out the SFP™ code on that NNP®, and are developing a demonstration capability on that platform.
correlating the two signals.

6. Summary, Conclusions, and Recommendations

We have developed a novel sensor fusion capability, modeled after investigations of sensor fusion in the superior colliculus of mammals by our colleagues at Wake Forest University. We have studied and evaluated the performance of this new capability, and found that it is useful for extracting weak targets in a noisy background. This method results in a 5 – 7 dB improvement in target detection as compared to the standard single-sensor method. This method has been patented (patent number 5850625).
We have designed a Sensor Fusion Processor that will provide a unique hardware embodiment of this sensor fusion concept. The components include an Input/Output Processor (IOP), a PCI interface and the AAC NNP™. The SFP™ show great promise as a high-speed fusion processor. Both the SFP™ and its components can be marketed separately, and there has been commercial interest in both the whole unit and the components.
Recommendations: The biologically inspired sensor fusion concept should be exploited for a wide range of sensor fusion applications, ranging from radar and other multisensor target detections to sonar, and to other areas in which a topographic representation of contact information is appropriate. This will include many forms of imagery. Immediate applications should include such things as fusing single or multiple radar returns with IFF. Although one of these fusion efforts is currently available, we believe that better target detection/fusion could be made with our method, providing improved target detections while still maintaining relatively low false alarms.

The concepts developed under this effort should be further developed to take into account the most recent results of biological studies. This should include the influence of context, of variable shape for the receptive fields and sensor fusion functions, and other factors..

The SFP™ prototype that AAC has been developing should be tested in a variety of real-world applications, with a view to transitioning this new hardware capability into the fleet.

7. Acknowledgments

The authors will to thank our fellow researchers at Accurate Automation Corp as well as the research team led by Prof. Barry Stein at the Bowman-Gray School of Medicine, Wake Forest University for their research in the biological aspects of sensor fusion. Special thanks are due to Mr. John Baron, DoD SBIR Program Director, Messrs. Vincent Schaper and John Williams, ONR SBIR Program and Deputy Program Manager, respectively. All research described in this paper has been supported by the Small Business Innovation Research (SBIR) program and the Small Business Technology Transfer Research (STTR) program.

8. References

The Department of Defense, Joint Requirements Oversight Council (199**7**), *Joint Warfighting Science and Technology Plan*

The Department of the Navy (August 1998). *Science and Technology Requirements Guidance (CNO-N091.*

The Department of the Navy (1997). *Technology Issues from the Fleet Commander-in-Chiefs* (STRG 1996*).*

Granger, R., Ambros-Ingerson, J., Anton, P.S., Whitson, J., & Lynch, G. (1992). "Computational action and interaction of brain networks," in S.F. Zornetzer, J.L. Davis, & C. Lau (Eds.), *An Introduction to Neural and Electronic Networks*. New York: Academic.

McHaffie, J.G., Norita, M., Dunning, D.D., and Stein, Barry E. (1993). "Corticotectal relationships: direct and "indirect" corticotectal," *Progress in Brain Research*, Vol. 95, Elsevier Science Publishers.

Meredith, M.A. and Stein, B.E. (1986a). "Visual, auditory and somatosensory convergence on cells in superior colliculus results in multisensory integration," *J. Neurophysiol*, 56: 640-662.

Meredith, M.A. and Stein, B.E. (1986b). "Spatial factors determine the activity of multisensory neurons in cat superior colliculus, *Brain Res.*, 365: 350-354.

Meredith, M.A., Nemitz, J.W. and Stein, B.E. (1987). "Determinants of multisensory integration in superior colliculus neurons," I. Temporal Factors. *J. Neurosci.*, 10:3727-3742.

Niida, T., Stein, B.E., & McHaffie, J.G. (1997). "Response properties of corticotectal and corticostriatal neurons in the posterior lateral suprasylvian cortex of the cat*," J. Neuroscience*, 1**7** (21, November 1, 1997), 8550-8565.

Stein, B.E., Meredith, M.A., and Wallace, M.T. (1993). "The visually responsive neuron and beyond: multisensory integration in cat and monkey," *Progress in Brain Research*, Vol. 95, 79-89.

Stein, B.E., & Meredith, M.A. (1993). *The Merging of the Senses*. Cambridge, MA: MIT.

Stein, B.E., & Meredith, M.A. (1990). "Multisensory integration: Neural and behavioral solutions for dealing with stimuli from different sensory modalities." In A. Diamond (Ed.), *The Development and Neural Bases of Higher Cognitive Functions, Annals of the New York Acad. of Sciences, V. 608*, 51-65.

Wallace, M.T., & Stein, B.E. (1994). "Cross-modal synthesis in the midbrain depends on input from cortex," *J. Neurophysiology*, **71** (1, January, 1994), 429-432.

Welch, R.B. and Warren, D.H. (1986). "Intersensory interactions," In: K.R. Boff, L. Kaufman and J.P. Thomas (Eds.), *Handbook of Perception and Human Performance, Vol.I: Sensory Process in Perception*, Wiley, New York, pp. 1-26.

Wilkinson, L.K., Meredith, M.A., & Stein, B.E. (1996). "The role of anterior ectosylvian cortex in cross-modality orientation and approach behavior," *Exp. Brain Research*, **112**, 1-10.

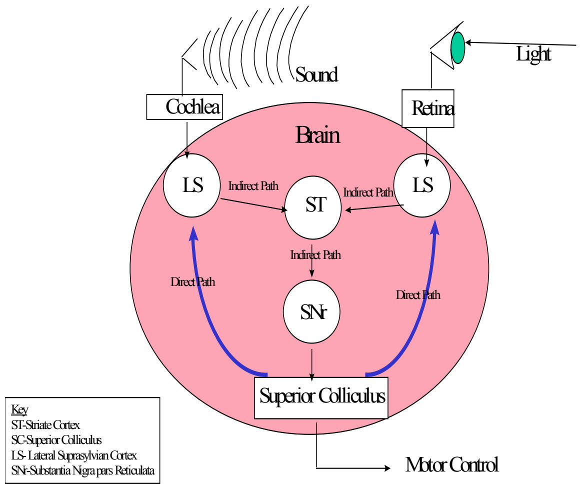

Figure 1: Organization of direct and indirect pathways for sensor fusion in the SC.

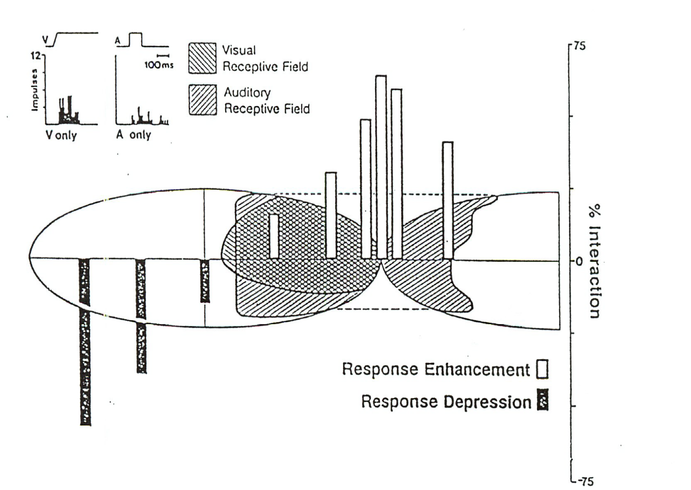

Figure 2 Biological Sensor Fusion Across Loose Spatial Registry (Stein & Meridith, 1993)

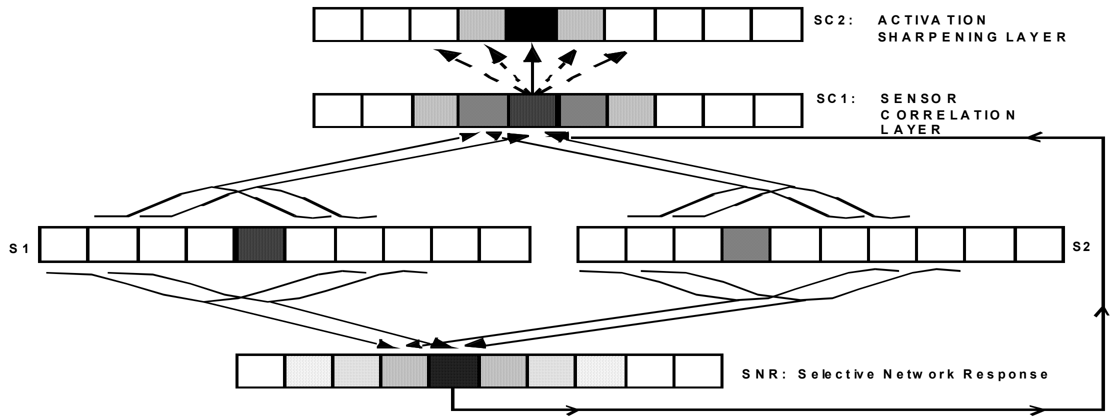

Figure 3. Basic SC-inspired sensor fusion architecture

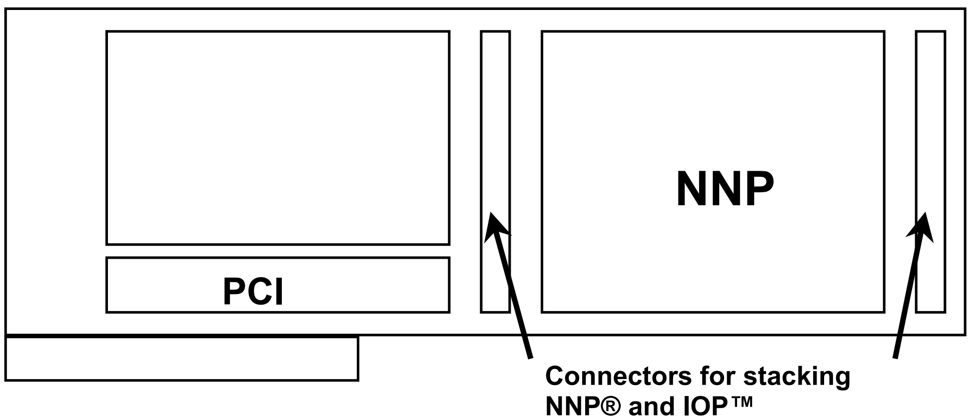

Figure 4. Initial design for SFP™

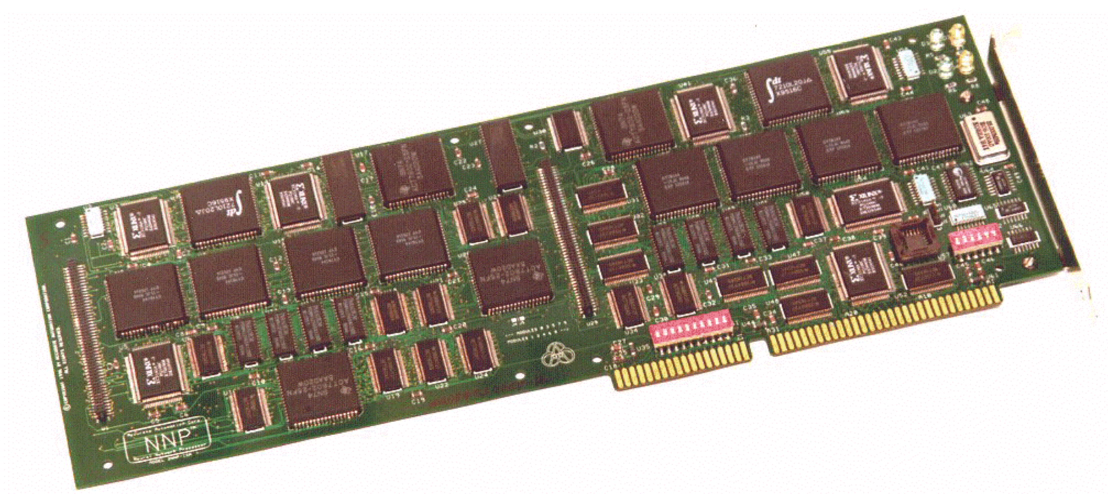

Figure 5. Dual-processor NNP configuration